Challenge: Towards Distributed RFID Sensing with Software-Defined Radio

Danilo De Donno
University of Salento
Via per Monteroni - Ecotekne
73100 Lecce, Italy
danilo.dedonno@unisalento.it

Fabio Ricciato
University of Salento
Via per Monteroni - Ecotekne
73100 Lecce, Italy
fabio.ricciato@unisalento.it

Luca Catarinucci
University of Salento
Via per Monteroni - Ecotekne
73100 Lecce, Italy
luca.catarinucci@unisalento.it

Angelo Coluccia
University of Salento
Via per Monteroni - Ecotekne
73100 Lecce, Italy
angelo.coluccia@unisalento.it

Luciano Tarricone
University of Salento
Via per Monteroni - Ecotekne
73100 Lecce, Italy
luciano.tarricone@unisalento.it

ABSTRACT

Current Radio-Frequency Identification (RFID) technology involves two types of physical devices: tags and reader. The reader combines in a single physical device transmission (to the tags) and reception (from the tags) functions. In this paper we discuss an alternative approach, where receive functions are performed by a separate device called "RFID listener". This allows distributed tag-sensing schemes where one transmitter coexists with multiple listeners. We discuss pros and cons of both approaches and present our implementation of a passive RFID listener on GNU Radio. Our implementation is a basis for experimenting with future distributed listener-based systems, but it can be also used as a cheap and flexible protocol analyzer for currently available commercial RFID readers.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]

General Terms

Measurement, Performance, Design, Experimentation.

Keywords

RFID, Software-Defined Radio, GNU-Radio, WSN.

1. INTRODUCTION

Radio-Frequency Identification (RFID) technology is being adopted in a growing range of application fields, from logistic to inventory management. Standardization has not yet converged to a unique worldwide solution, and several competing standards exist nowadays. This study focuses on purely passive tags, operating in the UHF range 865.6÷867.6 MHz with EPC Class-1 Generation-2 (Gen2) standard [1]. Regardless of standard issues, all current RFID systems involve two types of components: tags and readers. For purely passive tags, the reading process forsees three logical functions:

- transmit the queries to the tags;
- transmit a Current Wave (*CW)* to energize the tags;
- receive and decode the replies from the tags.

In traditional RFID technology all such functions are concentrated at a single physical device, i.e. the reader (ref. Figure 1a). Here we envision an alternative approach where they are distributed among multiple devices. More specifically, we decouple the transmission functions (*Query* and *CW*) from the reception, and consider a scheme where transmitter and receiver are physically separated (ref. Figure 1b). We refer to the receiving device as "RFID listener". It should be remarked that, in order to decode the messages from the tags, the listener must decode the signals from the transmitter as well: this allows to decode the query and to identify the time/frequency of the tag reply.

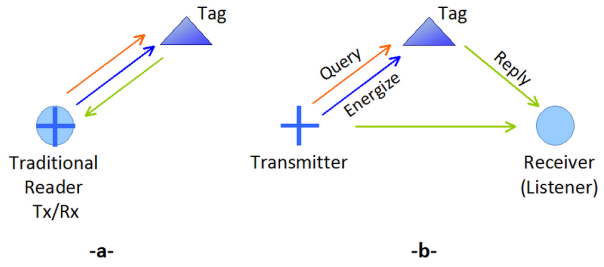

Figure 1. Conventional Reader (a) and separate Transmitter/Listener (b)

The separation of transmission and reception functions enables novel operational schemes where a single RFID transmitter coexists with multiple RFID listeners, and all such devices

MobiCom'10, September 20-24, 2010, Chicago, Illinois, USA.

cooperate to discover, read and localize the tags in a distributed fashion. In such a vision, RFID listeners can be seen as "tag sensors". They could be integrated into the low-cost low-power nodes of a Wireless Sensor Network (WSN) operating e.g. with IEEE 802.15.4 radios. Reception by multiple listeners brings multiple advantages over the multi-reader scenario, especially in case of dense deployments, which we highlight in Section 2.

Note that the Gen2 protocol (summarized later in Section 4.2) requires a tight coordination between the Transmitter and Receiver, mainly to support explicit acknowledgments. Therefore, transmitter and listener cannot act independently as far as the Gen2 protocol is used. The distributed scheme envisioned here, where independence between transmitter and listener is a key feature, requires the definition of alternative protocols. This paper is not meant to propose any particular solution, but rather to motivate further research in this direction.

In order to carry out experimental research in the new scenario and test different protocol options, it is desirable to have a flexible implementation of both transmitter and listener(s). Software-Defined Radio (SDR) [5], and particularly the open-source platform GNU-Radio [10], offers a cheap and versatile solution to researchers and practitioners to setup up a test-bed environment.

In this paper we make three main contributions. First, we discuss the potential pros and cons of the envisioned distributed scenario, where RFID transmitters and listeners are physically separated. Second, we present the initial implementation of a RFID listener based on GNU-Radio, which will be made freely available on CGRAN archive [6]. Our implementation is fully complaint with Gen2, and represents a starting basis for developing novel versions of listeners going beyond the current specifications, for experimental pre-standard research. Since our RFID listener is able to decode both the queries from the reader and the replies from the tag(s), it can be also used as a sort of passive "protocol analyzer" for conventional Gen2 systems. As a third contribution, we present some measurement results from a commercial RFID system to illustrate the use of the listener as a protocol analyzer.

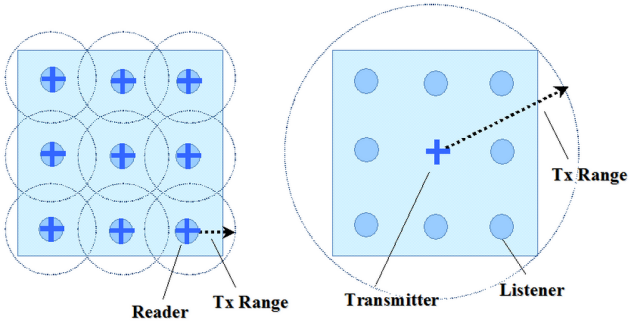

Figure 2. Traditional reader-based scenario (left) and 1:n cooperative listener-based scenario (right).

The idea of leveraging GNU-Radio for the analysis of conventional RFID communications was pioneered by Buettner and Wetherall in [7]. They used a very preliminary GNU-Radio implementation to decode the signals from the reader on a fixed frequency. Therefore the protocol analysis in [7] had to rely on the logs from the reader under test. The same authors have recently released a GNU-Radio implementation of a complete reader (see [8] and [9]) following the conventional approach where a single device acts as transmitter and receiver. As such, the Buettner's reader is able to decode only the tags it queries directly, while our implementation can "opportunistically" decode tag signals queried by external devices.

2. READER-BASED VS. LISTENER-BASED DISTRIBUTED READING

In practical RFID systems limitations exist both on the maximum *reception range* from the tag – denoted by R_L – and on the range of *CW* energization range – denoted by R_T. The former derives from the limit on the maximum power *scattered* by the tag, while the latter depends on the maximum power emitted by the transmitter. We are interested in cases where the energization range can be made greater than the reception range, i.e. $R_T > R_L$. While in the traditional scheme the reading range equals the minimum of the two, i.e. $R_R = min(R_L, R_T)$, as a single physical device performs both transmission and reception, in the listener-based scheme the two ranges are decoupled. In principle, this allows a single higher-power transmitter to illuminate the whole area of interest while multiple listeners are deployed to ensure full sensing, as shown in Figure 2b. For simplicity we consider the static case where all infrastructure devices are fixed.

In this section we compare two alternative system approaches: the "traditional" scheme based on the use of current RFID readers (Figure 2a), and the cooperative listener-based scheme where a multiple Listeners cooperate to sense the area illuminated by a single Transmitter (Figure 2b) – denoting by n the number of listeners, we will refer to the latter as "1:n listener-based" scenario. In principle one can consider more elaborated settings with multiple transmitters (m:n, with $m << n$ the number of transmitters) in order to extend the system coverage beyond the area illuminated by a single Transmitter. In such cases, all the arguments introduced hereafter in support of the 1:n model still apply within a single illuminated area.

Our aim is NOT to claim superiority of the 1:n listener-based approach over the traditional reader-based one: both approaches have different pros and cons, and depending on the particular application one scheme might be more or less convenient than the other. However we will show that the 1:n listener-based scheme becomes particularly attractive in dense scenarios with many short-range infrastructure devices. In fact, using passive listeners in place of active readers brings in a number of advantages in terms of deployment cost, power consumption, capacity, sensing efficiency and localization accuracy that reinforce each other and become important in dense scenarios with many coexisting devices. Therefore, we believe that further research on the cooperative scheme with transmitter/listeners separation is worth to be undertaken to unveil its potential, particularly for dense deployment scenarios and/or RFID/WSN integration.

Cost, size and power gains. Compared to traditional readers, which include also transmission functions, RFID listeners are receiver-only UHF radios. They save the whole transmission chain, including important components like the Power Amplifier and Digital-to-Analog Converter. This results into three important device-level savings, which reinforce somehow each other: lower fabrication cost, smaller size and much lower power consumption. The latter in turns prolongs the battery lifetime. Furthermore it opens up the possibility of building self-alimented devices

coupled with energy-harvesting modules. The low-cost low-power nature of such devices facilitates the integration in WSN nodes, as discussed below.

Enabling Denser Deployment. The savings mentioned above – fabrication cost, size and power – concur to lower the cost of infrastructure deployment and maintenance, the latter being heavily influenced by battery replacement. In turn, such gain can be used to *increase the density* of deployed listener devices. In other words, it becomes cost-effective to deploy a higher number of receive-only listeners to cover a given area of interest. Higher-density deployment brings two major benefits. First, it increases the gain of the cooperative reception schemes discussed below, thus increasing the system capacity in terms of spatial and/or temporal reading rate. Second, it allows for better accuracy in the localization of fixed tags and/or in the tracking of moving tags: assuming that multiple listeners placed at known locations "hear" the same tag, well-known localization techniques based on the received signal strength (see e.g. [21]) can be applied to improve the localization resolution within the proximity range of listeners.

Reduced transmission coordination. To illustrate this point, we first note that in a traditional multi-reader system two kinds of collisions occur:

- Reader Collisions: between CW/Query signals from different readers hitting the same tag;
- Tag Collisions: between replies from different tags activated by the same reader.

Tag collisions can be partially recovered in the listener-based scheme via cooperative reception, as discussed later. Instead, reader collisions are always destructive events: if queries from different readers reach the tag at the same time, the tag becomes unable to issue a valid reply even if they were transmitted on different frequency channels. Consequently, if multiple readers are present in the same environment, they need to implement some sort of Medium Access Control (MAC) mechanism to prevent reader collisions in time (see e.g. [20]). Instead, in the 1:*n* listener-based scheme there is only a single active Transmitter, which removes completely the need for transmission coordination among infrastructure devices. In general, in *m*:*n* listener-based scenarios the number of active devices is reduced from *n* to *m*<<*n*, thus lowering the coordination overhead among transmitters.

Saving on the MAC reduces the *complexity* of the whole system and improves its *scalability*. Such system-level gains, coupled with the device-level gains discussed above, contribute further to the feasibility and cost-effectiveness of high-density deployments with many listeners.

Temporal efficiency. The adoption of a unique transmitter in 1:*n* systems increases the temporal efficiency as it avoids spurious reader collisions (for collision-oriented MAC) or time-division (for collision-free MAC). Refer to Figure 3 and assume that Tag-A is in the reception range of device L1, while Tag-B is in the reception range of L1, L2 and L3. If we used traditional Readers instead of Listeners, we should implement a *K*-slot time-division scheme to separate the transmissions from different readers (e.g. *K*=3 in the case of Figure 3). Therefore, reading opportunities for Tag-A occur only *1/K* of the time. Instead, with 1:*n* listener-based scheme, each tag can be queried at every slot.

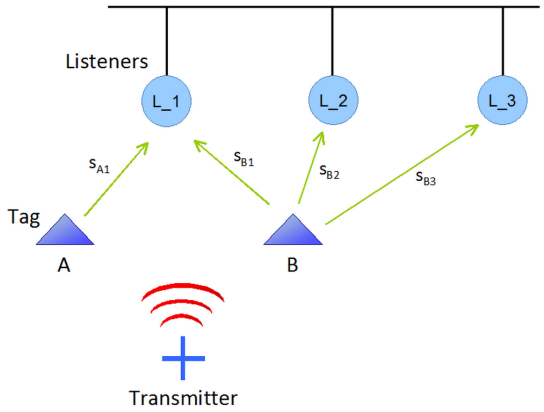

Figure 3. Example of cooperative reception.

Cooperative Reception. The 1:*n* listener-based scheme paves the way for implementing well-known cooperative reception techniques like Soft Combining and Interference Cancelation. To illustrate these concepts, refer again to Figure 3. Denote by $s_{X,k}$ the frame from Tag-*X* as received from Listener *k*. Assume that Tag-B's signals are received incorrectly by listeners L2 and L3 due to low SNR, i.e. both $s_{B,2}$ and $s_{B,3}$ frames fail the CRC. Furthermore assume that $s_{A,1}$ and $s_{B,1}$ collide on listener L1. In such a scenario, no message can be correctly decoded as long as each receiver operates in isolation, as in the traditional reader-based scenario. On the contrary, by adopting cooperation among passive listeners, the reply from both tags can be correctly decoded. First, Chase combining [17] can be applied to the signals $s_{B,2}$ and $s_{B,3}$ so that the combined SNR becomes sufficient to correctly decode the Tag-B frame. Chase combining is a simple form of Soft Combining: it can be applied before demodulation on down-converted samples, or after demodulation at the bit level – the former variant requiring considerably less communication bandwidth in backplane at the price of lower gain.
Once that Tag-B message has been correctly recovered, its re-modulated version can be subtracted from the signal $s_{A,1}+s_{B,1}$ received by L1: in this way, the "known" interference $s_{B,1}$ is removed and the residual signal can be de-modulated to recover $s_{A,1}$. This is a simple form of Interference Cancellation (see e.g. [18] and references therein).

Soft Combining and Interference Cancellation are among the most basic ingredients of multi-user detection. In fact, at a more abstract level, we can consider a group of *h* active tags – queried and energized by a single Transmitter – and a set of neighboring *n* listeners as a simple form of *h*x*n* multi-user single-input multiple-output (MU-SIMO) system. The application to RFID of multi-user detection has been proposed recently to increase the capacity of traditional readers in single-device scenarios – see [19] and references therein – where the device size imposes a severe limit on the number and spacing of antennas, hence on the multiplexing gain. Instead, in the multi-device version envisioned here it is possible to achieve much higher reliability and throughput via increased diversity, also in case of harsh propagation environments where the tag-listener channel might be obstructed (e.g. by metal or liquid materials) and/or disturbed by external sources of interference. More in general, we believe that the perspective of considering a distributed scenario with many passive listeners and tags acting as a sort of large multi-user

system, opens up interesting directions for building high-capacity, more robust RFID sensing systems with built-in localization and tracking features. A pre-requisite for such vision is to make the tag reception independent from the query/energization functions.

3. POTENTIAL APPLICATIONS

A key system-level aspect concerns the inter-device communication channel: regardless on whether we adopt traditional Readers or Listeners, such devices need to communicate to exchange and/or report the readings to a central sink, therefore a separate control channel is required in any case. We envision two reference scenarios that might be of interest in different applications.

One possibility is to connect the infrastructure devices to a wired backplane (as in Figure 3): in this case the higher bandwidth available for inter-device communication could support the more sophisticated cooperative reception schemes in the listener-based scenario, e.g. pre-demodulation soft-combining, thus achieving higher cooperation gain at the cost of heavier infrastructure deployment.

At the other extreme, one can resort to a wireless backplane channel like the cheap IEEE 802.15.4, nowadays the most popular choice for WSN. Several previous works have considered RFID/WSN interworking in the framework of conventional reader-based scheme (see e.g. [2,3]). In the cooperative listener-based scenario the integration can be even tighter: the RFID listener can be seen as purely passive "tag sensing" module that in principle can be integrated onboard of a low-cost low-power WSN node along with other sensors. In this sense one can see the WSN "augmented" with RFID sensing capabilities – provided that one independent RFID transmitter illuminates the environment.

In both cases – wired and wireless backplane – the infrastructure cost is expected to be much lower for the listener-based scheme compared to the traditional reader-based solution, thus making high-density deployments with many infrastructure devices more cost-effective. One might think to several scenarios where the RFID sensing infrastructure is "diffused" in a building, e.g. think to small listeners deployed every few meters, below the floor tiles or embedded in a smart sticky tape attached to walls, ceiling or along corridor paths.

Such "diffused infrastructure" scenarios are attractive for those real-world applications that can directly benefit from denser listener deployment. Recalling the discussion in Section 2, this is the case where one or more of the following requirements apply:

1. accurate localization and tracking of moving tags are needed (beyond tag identification);
2. high reading capacity is required, i.e. many tags must be read at the same time in a given area.;
3. the operating environment is harsh for radio signals, e.g. rich of metal obstructions and/or external sources of interference.

Applications posing all three types of requirements are easily found in production and logistics: factories and warehouses are typical examples of harsh radio environments traversed by many units of goods that need to be identified, localized and tracked continuously. Also, dense RFID systems might find applications in large shopping centers and malls, for those goods that need to be searched or tracked across their entire path within the shop, rather than at specific pre-determined spots like exit gates.

4. RFID LISTENER IN GNU-RADIO

4.1 Motivations

We developed a Software-Defined Radio (SDR) receiver on the GNU-Radio platform [10] that, located near a commercial RFID deployment, is able to monitor and decode the transmissions from the reader and from the tags. Our receiver provides a basis for future prototypes of RFID listeners, to be used in experimental pre-standard research for the cooperative 1:n scenario discussed above. It can also serve as a complete protocol analyzer for commercial Gen2-compliant RFID systems – note that existing commercial readers return only high-level data and logs but do not provide the means to observe completely what happens at MAC and PHY layers.

Our implementation is based on the Universal Software Radio Peripheral (USRP) hardware [11] and the GNU-Radio toolkit. The USRP is a general purpose RF front-end for SDR development, connected via USB 2.0 with a standard PC where the signal processing is performed by the GNU-Radio toolkit. The platform has some well-known limitations. First, the USB 2.0 connection limits the maximum input signal bandwidth to approximately 8 MHz. This is not a problem for our implementation since the European Gen2 specifications at 866.5 MHz are limited to 2 MHz, well within the limit. Second, the software-based processing introduces high latency and consequently inflates the Tx/Rx turnaround time. This is not a serious problem for our receive-only listener, since the Tx chain is not used. A third and more serious limitation refers to the limited ADC dynamic range. The problem is originated from the contemporary transmission of large-amplitude *CW* (by the reader) and low-amplitude replies by the tag(s). When the amplitude gap is large, tuning the dynamic ADC range to the *CW* amplitude, to avoid non-linear distortions, results in the tag signals being lost in the quantization error. Commercial readers generally use narrow band-pass filters *before* ADC conversion of the up-link signal [12], instead USRP does not provide any type of signal processing before digital conversion. In practice, this problem imposes a limitation on the maximum sensing range of RFID listener in USRP, which is evaluated later in Section 4.

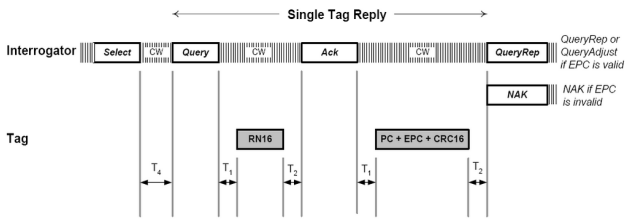

Figure 4. MAC for Gen2 protocol.

4.2 Overview

The modulation scheme used by Gen2 for up-link and downlink transmission is Amplitude Shift Keying (ASK) and the MAC protocol is based on Framed Slotted Aloha [4]. Figure 4 shows the sequence of commands that constitute an *Inventory Round.* When a tag receives a *Query* command, it picks a random 16-bit number (*RN16*) in [0, $2Q - 1$], with $0 \leq Q \leq 15$ a configurable parameter,

and stores it in the slot counter which is then decremented at each *Inventory Round*. When the counter gets to 0, the tag transmits its *RN16*. Upon reception, the reader will echo the *RN16* in the *ACK* message. If the tag successful receives the *ACK* with the correct *RN16* number, it will finally backscatter its 128-bits ID (*EPC* message in Figure 4).

4.3 Implementation

Single-antenna version. The block diagram of the software sub-system for the single-antenna version is shown in Figure 5. In the upper branch the first block is a matched filter configured to maximize the signal-to-noise ratio (SNR) of reader and tags transmissions. The matched filter acts also as a tunable band-pass filter to select the channel of interest – its center frequency can be varied by tuning the phases of the complex taps. The following ASK demodulation block transforms a stream of complex I/Q values into a stream of amplitude values. The resulting stream contains both the reader commands and tag transmissions. According to the "*reader-talks-first*" paradigm of Gen2, a tag response can occur only after a reader command which provides to the tag a number of parameters that shape its backscatter signals. Therefore in order to identify and decode the tag signal one must first decode the reader command.

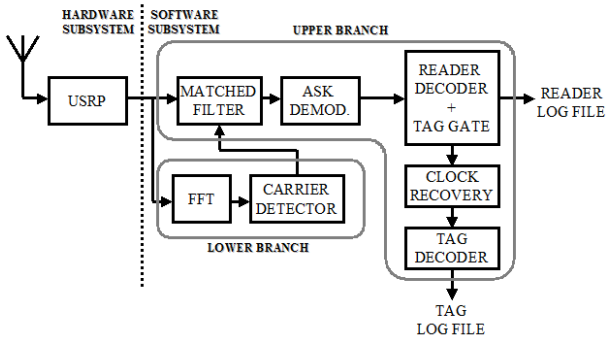

Figure 5. Block diagram of the 1-antenna receiver.

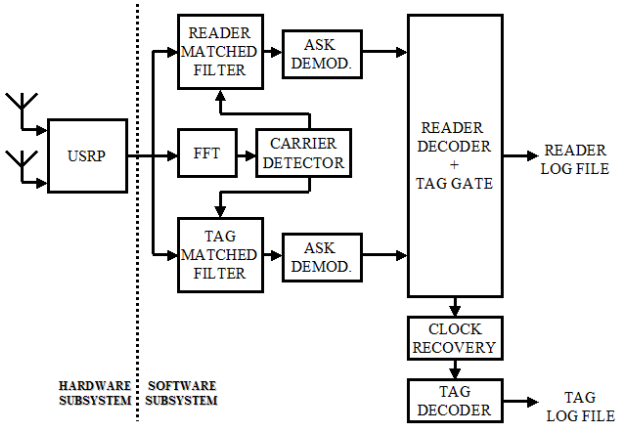

Figure 6. Block diagram of the 2-antennas receiver.

Recall that the RFID listener must decode two types of transmissions: from the reader (e.g. *ACK, Query, CW*) and from the tags (*RN16, EPC*). We developed two different versions of the listener: with single and double antenna. The block diagrams are given in Figure 5 and Figure 6 respectively. In the dual-antenna version, one antenna is dedicated to receive the signals from the reader, while the other collects the low-power signals from the tags. In principle, the two antennas could be directional and oriented towards different directions. Further evolutions of the system could foresee adaptive antenna arrays.

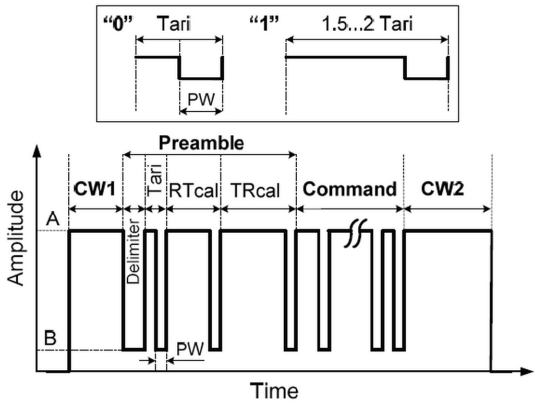

Figure 7. PIE symbols (top) and baseband reader transmission (bottom).

In Gen2 Pulse Interval Encoding (PIE) is used as downlink encoding scheme: the '1' and '0' bits are encoded by pulsing the carrier wave *CW* at different time intervals, as sketched at the top of Figure 7. A typical Gen2 reader transmission is shown at the bottom of Figure 7. The reader transmits a un-modulated carrier wave (*CW*) followed by the modulated portion and a second *CW* signal (*CW2*). The modulated signal consists of the *preamble* – which includes several critical time periods that define link timing between the reader and the tag – and of a *command* to the tag. The parameters of the time intervals used for the preamble detection of the reader-decoding block are summarized in Table 1.

Table 1. GEN2 Protocol parameters

Name	Allowed values	Explanation
CW1	Unspecified	Power-Up
Delimiter	12.5 us +/- 5%	Start of frame
Tari	6.25us .. 25us	'0' duration
PW	0.265Tari .. 0.525Tari	Pulsewidth
RTcal	2.5Tari .. 3Tari	'0'+'1' duration
TRcal	1.1RTcal .. 3RTcal	U.L. calibration
Command	Binary sequence	Command
CW2	Unspecified	Tag Response

Once the last pulse of a reader command is detected, a tag backscatter can occur so we un-gate the signal and pass it to the next "Clock Recovery" block which re-samples and interpolates the data stream. The now "clean" signal of tag backscatter is passed to the "Tag Decoder" block. In order to detect the preamble and correctly decode the tag's message, the reader's message must have been preliminarily decoded. In fact, in Gen2 the reader chooses the up-link parameters and communicates them to the tags in the opening symbols of each packet. For these reason, decoding reader commands is required to access the up-link parameters and best tune-up in *real-time* both the matched filter and the "Tag Decoder" block (ref. Figure 8). In particular

the preamble of the reader commands contains the tag-to-reader calibration symbol *TRcal* (see Table I) and the divide ratio (*DR*) parameter. The tag can then determine its backscatter link frequency (*BLF*) as *BLF=DR/ TRcal*.. Our listener determines the BLF from the reader command and passes it to the matched filter block, which adapts its center frequency accordingly.

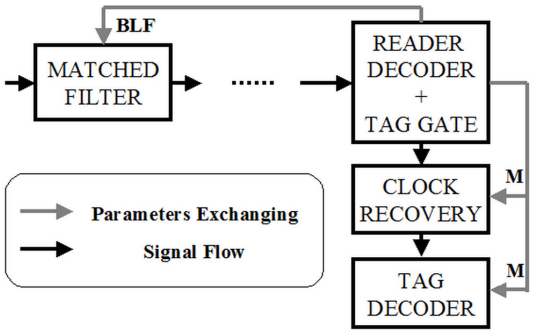

Figure 8. Exchange of parameters between blocks.

The tag shall encode the backscattered data as either FM0 baseband signaling or Miller-modulated subcarrier (MMS) encoding. In FM0 a binary '0' has a transition in the middle of a symbol, whereas a binary '1' does not. In MMS the FM0 signal is multiplied by a square wave with either *M*=2, 4 or 8 periods for each FM0 symbol. The reader sends the *M* parameter (*M*=1 maps to FM0) in the same command along with *TRcal* and *DR*. From it the tag can determine the encoding scheme and the related data-rate *BLF*/*M*. Our listener extracts the latter from the reader command and communicates it to the "Clock Recovery" and "Tag Decoder" blocks (ref. Figure 8).

The "Tag Decoder" block detects the preamble via correlation and decodes the subsequent bits using also a correlator. Recall that tag can backscatter two possible messages: a *RN16* message as response to a reader *Query* command *(or QRep)*, or a *EPC* message following a correct *ACK* command. *EPC* message include the "Protocol Counter" (*PC*) field and a *CRC-16* (bit) to check if the backscattered ID has been successfully decoded.

Note that the "Reader Decoder + Tag Gate" and "Tag Decoder" blocks are similar to the homologous blocks used by Buettner in [7-9] but several parts were modified to adapt to our receiver-only scenario and some novel features were added: "Tag Gate" functionality, parameters exchange with other blocks (see Figure 8) and RSS estimation.

Dual-antenna version. The block diagram of the dual-antenna implementation is shown in Figure 6. Compared to the previous version the main difference lies in the overall architecture rather than in the functionalities of the individual blocks. For this reason we only present the main features that distinguish it from the single-antenna version.

The USRP delivers to the host PC an interleaved signal containing alternating I/Q samples from the two antennas. So the first block of the software sub-system (omitted in Figure 6 for simplicity) is a de-interleaver that separates the two streams and sends them to the corresponding receiver chain. Now each signal first pass through its matched filter and then through the ASK demodulator. The "Reader Decoder + Tag Gate" block is quite different from that used for the single-antenna receiver (see Figure 9).

The single-antenna version of the block has only one input, so the gated/un-gated signal is the same entering the reader command decoder. The dual-antenna version instead receives two inputs, therefore in order to gate/un-gate the tag stream the two receive chains must be perfectly synchronized. To this end the first signal must be delayed in order to compensate for the offset due to a different number of taps in the two matched filters. Afterwards it is decimated/interpolated to fit the data-rate of the reference signal. When un-gated, the signal containing a tag response is passed to the "Clock Recovery" and "Tag Decoder" blocks which are similar to the single-antenna version.

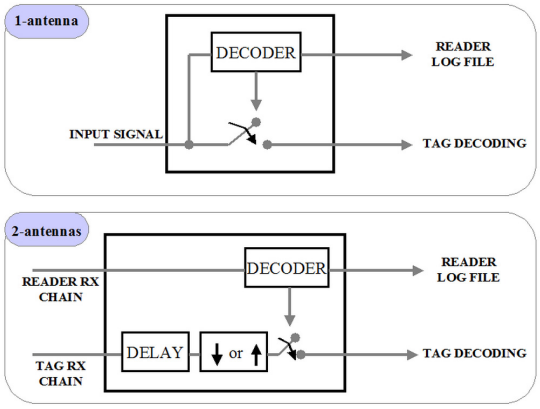

Figure 9. "Reader Decoder + Tag Gate" block in case of 1-antenna (top) and 2-antenna (bottom).

Frequency tracking sub-system. According to the international regulations RFID readers are allowed to use a channel for a maximum period of time that varies from 0.4 seconds, set by the U.S. *FCC* [13], to 4 seconds recommended by *ETSI* [14]. This means that commercial readers must implement frequency-hopping schemes with pre- defined or pseudo-random sequence. In order to decode reader and tags transmissions in the 2 MHz acquired band (865.6-867.6 MHz) it is necessary to adopt a "frequency tracking system". In particular, the center frequency of the matched filter should be adjusted to the instantaneous frequency used by the reader by changing the phase of the complex taps. Both versions of the listener use a frequency tracking system based on a 64 samples FFT calculation followed by a block that finds the peak reader carrier corresponding to the most powerful bin. The detected frequency is then communicated to the matched filter blocks (see the lower branch of Figure 5 and the central branch of Figure 6).

5. EXPERIMENTAL RESULTS

This section presents some preliminary experiments performed to test the correct operation of the listener implementation. The experiments were conducted in a standard office room of approximately 7x6x4.5 meters. The RFID deployment consists of one Alien ALR-8800 reader equipped with two Alien ALR-9610 circular antennas, and eight Alien 9640 "Squiggle" tags. The tags were positioned on a poster board in a 2x4 grid with approximately 10 cm spacing. The reader, tags' panel and listener

were arranged on the horizontal plane at 2 meters height as shown in Figure 10.

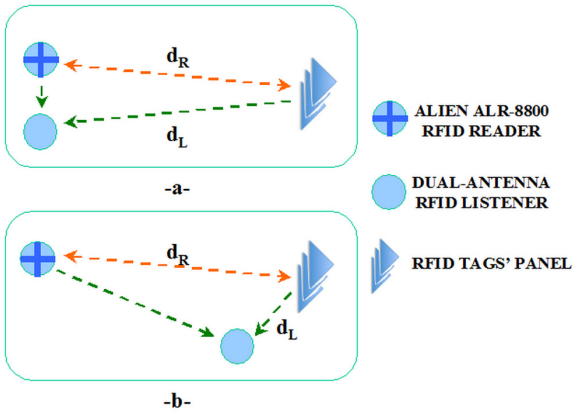

Figure 10. Topologies used for the experiments.

Two different sets of experiments were conducted: in EXP1 (see Figure 10a) the listener was located near the reader ($d=d_R=d_L$) while in EXP2 (see Figure 10b) the listener was kept at fixed distance d_L=1 m from the tag. For each experiment we conducted several measurements at different reader-tag distance $d_R=d$. The dual-antenna version of the RFID listener was used on a general purpose PC (Intel Core2 Quad CPU @ 2.83 GHz, 2 GB RAM) with Ubuntu 9.10 32-bit and GNU-Radio 3.2.2 toolkit installed. The USRP motherboard was equipped with two RFX-900 daughter-boards and Alien ALR-9610 circular antennas.

For each experiment we let the reader and listener active for 10 minutes and record the following counters:

- N_Q = # of queries issued by the reader
- $N_{R,OK}$ = # of successful *EPC* readings by the reader
- $N_{R,ERR,v\text{-}ACK}$ = # of erroneous *EPC* readings by the reader with a corresponding valid ACK
- $N_{R,ERR,i\text{-}ACK}$ = # of erroneous *EPC* readings by the reader with a corresponding invalid ACK
- $N_{L,OK}$ = # of successful *EPC* readings by the listener
- $N_{L,ERR,CRC}$ = # of erroneous *EPC* readings by the listener due to CRC errors.
- $N_{L,ERR,MISS}$ = # of *EPC* messages missed by the listener.

It holds that $N_{L,OK} + N_{L,ERR,CRC} + N_{L,ERR,MISS} = N_Q$. Note that all such quantities were measured by the listener in real-time. We checked offline that the value of $N_{R,OK}$ reported by the listener was always equal to the number of readings reported in the logs of the reader under test.

Commercial RFID readers like the Alien ALR-8800 are reported to maintain a high rate of successful readings up to a distance of approximately 10 meters under ideal conditions [12]. Instead, the sensing range of our current listener implementation is limited to few meters, due to the saturation of ADC dynamic range by *CW*, a problem discussed before in Section 3.1. In the first experiment, EXP1, we aim at measuring the sensing range of the listener. We measure the Reader Success Ratio (R-SR), defined as $N_{R,OK}/N_Q$, and the Listener Success Ratio (L-SR), defined as $N_{L,OK}/N_Q$. We also calculate separately the percentage of *EPC* messages decoded by the listener due to CRC failure ($N_{L,ERR,CRC}/N_Q$) and the percentage of missed *EPC* messages ($N_{L,ERR,MISS}/N_Q$).

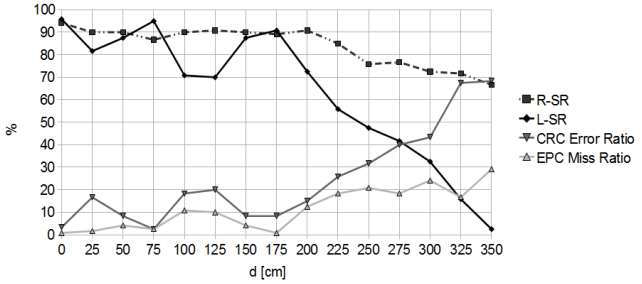

Figure 11. Performance achieved for the RFID listener.

Figure 11 summarizes these results. Up to c.ca. 2 meters both our listener and the commercial reader have comparable success ratios above 70%, with spatial fluctuations that can be accounted to the effect of multipath. Beyond 2 meters the reading rate of the listeners degrades steadily due to the ADC range limitation, and becomes practically zero at 3.5 meters. Interestingly the CRC error ratio increases more rapidly than the *EPC* miss ratio, which remains bounded to 20%-30%.

In the second set of experiments, EXP2, we study the performance of the Alien reader in terms of sensitivity at increasing of distance $d=d_R$; our RFID listener is held at d_L=1 m from the tags to minimize errors in decoding their messages. Figure 12 shows three metrics: the total *EPC* error ratio defined as $(N_Q\text{-}N_{R,OK})/N_Q$, the fraction of *EPC* errors with invalid *ACK*s sent by the reader ($N_{R,ERR,i\text{-}ACK}/N_Q$) – *ACK* is defined invalid when reader decodes an erroneous *RN16* message from the tag – and the fraction of *EPC* errors with valid *ACK*s ($N_{R,ERR,v\text{-}ACK}/N_Q$). Note that our listener is able to access such information since it can decode the transmissions from both the reader and the tags.

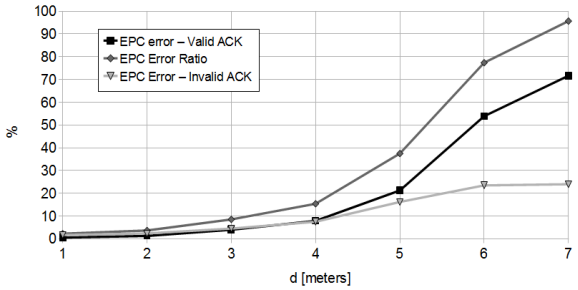

Figure 12. Gen2 reading performance.

It can be observed that at 7 meters approximately 70% of *EPC* errors had been correctly *ACK*ed. In other words, most of the times the reader was able to correctly decode the 16-bits string (*RN16*) but not the 128-bits ID from the same tag. This indicates that the use of shorter ID would lead to an increase in reading range and/or reading rate.

6. RELATED WORK

Previous works have studied various effects that degrade the performance of RFID systems but they lacked the instrumentation to determine the causes (see [15] and [16]). Regarding the SDR implementation, the closest work to our study is by Buettner and

Wetherall [22] which presents a preliminary version of a GNU-Radio receiver limited to decoding reader transmissions, while our implementation is able to decode the tag replies as well. Later the same authors developed a conventional RFID reader in GNU-Radio [8,9]. Our purpose is different, as our listener implementation is purely passive and is designed to decode tag/reader communications triggered by external interrogators. Another difference is that our implementation includes channel tracking, while the current version of the reader in [8,9] works on a fixed frequency.

The idea of combining RFID and WSN has received considerable attention recently by the research community. In most previous work, the basic idea is simply to leverage the WSN technology for coordinating multiple conventional readers to extend the coverage area (see e.g. [2,3]). The scenario considered here is different and foresees a tighter integration of the two radios: legacy WSN nodes are augmented with receive-only RFID listeners, while the interrogating and energizing functions are demanded to a separate transmitter acting as a RFID "lighthouse".

The application of MIMO to RFID was explored very recently by Langwieser *et al.* in [19] where a 2x2 single-device MIMO reader was presented. We consider here a fundamentally different scenario, where the physical separation between transmission and reception functions enables multi-device SIMO systems, with receiver macro-diversity across different listeners.

7. CONCLUSIONS AND ONGOING WORK

In this preliminary work we have drawn attention to an alternative approach to RFID reading, where transmit and receive functions to/from the tags are demanded to physically different devices. This enables "one-interrogates many-listen" scenarios, where a single tag transmission can be decoded in parallel by multiple listeners. These can cooperate to increase the spatial/temporal rate of successful tag readings. Furthermore, multiple reception can be leveraged to increase the tag localization accuracy. Being receive-only, the listener module is cheaper to produce and requires less power to operate: this makes dense deployments with large number of listener devices more cost-effective, which in turn helps to increase the cooperation gain, improves the localization accuracy, and facilitates the integration into the low-power low-cost nodes of a Wireless Sensor Network.

We have presented a first implementation of a Gen2-compliant RFID listener on GNU-Radio. As such, our listener can be readily used as a cheap and flexible protocol analyzer for conventional RFID systems. In the perspective of future research, it provides the basis for experimental versions of receive-only devices. Finally, our implementation can be useful for testing different protocols and prototyping future cooperative 1:n listener-based system.

8. REFERENCES

[1] EPCglobal. Epc radio-frequency identify protocols class-1 generation-2 uhf rfid protocol for communications at 860 MHz-960 MHz v. 1.0.9. 2005.

[2] Cheng Hsu et al. Enterprise Collaboration: On-Demand Information Exchange Using Enterprise Databases, Wireless Sensor Networks, and RFID Systems. *IEEE Trans. System, Man and Cybernetics, Part A*, vol. 37, pp. 519–532, 2007.

[3] J. Sung et al. The EPC sensor network for RFID and WSN integration infrastructure. *IEEE PerComW* 2007.

[4] L. G. Roberts. Aloha packet system with and without slots and capture. *SIGCOMM Comput. Commun. Rev.* 5(2) 1975.

[5] SDR Forum. SDR Forum Yearbook 2005. *SDR Forum Technical Conference*, Orange County, CA, 2005.

[6] The Comprehensive GNU Radio Archive Network (CGRAN). https://www.cgran.org.

[7] M. Buettner and D. Wetherall. An empirical study of UHF RFID Performance. *Proceedings of MobiCom'08*, 2008.

[8] Michael Buettner. Gen2 RFID Reader project website. https://www.cgran.org/wiki/Gen2.

[9] M. Buettner and D. Wetherall. A Flexible Software Radio Transceiver for UHF RFID Experimentation. UW TR: UW-CSE-09-10-02.

[10] GNU Radio. http://gnuradio.org.

[11] Ettus Research LLC. http://www.ettus.com/products.

[12] D. M. Dobkin. The RF in RFID: Passive UHF RFID in Practice. Elsevier. 2007.

[13] FCC regulations. PART 15, Radio Frequency Devices. http://www.fcc.gov.

[14] ETSI. EN 302 208 RFID standard. http://www.etsi.org.

[15] S. Aroor and D. Deavours. Evaluation of the state of passive UHF RFID: An experimental approach. In *IEEE Systems Journal*, vol. 1, pp. 168-176, 2007.

[16] K. Ramakrishnan and D. Deavours. Performance benchmarks for passive uhf rfid tags. *In Proceedings of the 13th GI/ITG Conference on Measurement, Modeling, and Evaluation of Computer and Communications Systems*, pp. 137–154, 2006.

[17] D. Chase. Code Combining – A Maximum-likelihood Decoding Approach for Combining an Arbitrary Number of Noisy Packets'. *IEEE Transactions on Communications*, vol. 33, pp. 385–393, May 1985.

[18] D. Halpering et al. Interference Cancellation: Better Receivers for a New Wireless MAC. *HOTNETS-IV*, Atlanta, USA, 14-15 November 2007.

[19] R. Langwieser, C. Angerer and A. Scholtz. A UHF Frontend for MIMO applications in RFID. *IEEE Radio and Wireless Symposium (RWS'10)*, New Orleans, January 2010.

[20] N. Vaidya and S. R. Das. RFID-based networks – exploiting diversity and redundancy. *ACM Mobile Computer and Communication Review*, 12(1), January 2008.

[21] A. Coluccia, F. Ricciato. On ML estimation for automatic RSS-based indoor localization. *IEEE International Symposium on Wireless Pervasive Computing (ISWPC'10)*, May 2010.

[22] M. Buettner and D. Wetherall. A "Gen 2" RFID monitor based on the USRP. *ACM SIGCOMM Computer Communication Review*, vol. 40, issue 3, pp. 41-47, July 2010.